

The Power of Teams that Disagree: Team Formation in Large Action Spaces

Leandro Soriano Marcolino¹, Haifeng Xu¹, Albert Xin Jiang¹, Milind Tambe¹, and
Emma Bowring²

¹ University of Southern California, Los Angeles, CA, 90089, USA
{sorianom, haifengx, jiangx, tambe}@usc.edu

² University of the Pacific, Stockton, CA, 95211, USA
ebowring@pacific.edu

Abstract. Recent work has shown that diverse teams can outperform a uniform team made of copies of the best agent. However, there are fundamental questions that were never asked before. When should we use diverse or uniform teams? How does the performance change as the action space or the teams get larger? Hence, we present a new model of diversity, where we prove that the performance of a diverse team improves as the size of the action space increases. Moreover, we show that the performance converges exponentially fast to the optimal one as we increase the number of agents. We present synthetic experiments that give further insights: even though a diverse team outperforms a uniform team when the size of the action space increases, the uniform team will eventually again play better than the diverse team for a large enough action space. We verify our predictions in a system of Go playing agents, where a diverse team improves in performance as the board size increases, and eventually overcomes a uniform team.¹

Keywords: Coordination & Collaboration · Distributed AI · Team Formation

1 Introduction

Team formation is crucial when deploying a multi-agent system [16, 7, 12, 15]. Many researchers emphasize the importance of diversity when forming teams [11, 10, 8, 14]. However, there are many important questions about diversity that were not asked before, and are not explored in such models. LiCalzi and Surucu (2012) [11] and Hong and Page (2004) [8] propose models where the agents know the utility of the solutions, and the team converges to the best solution found by one of its members. In complex problems the utility of solutions would not be available, and agents would have to resort to other methods, such as voting, to take a common decision. Lamberson and Page (2012) [10] study diversity in the context of forecasts, where the solutions are represented by real numbers and the team takes the average of the opinion of its members. Domains where the possible solutions are discrete, however, are not captured by such a model.

¹ This paper is the full version of our AAAI'2014 paper "Give a Hard Problem to a Diverse Team: Exploring Large Action Spaces", containing: (i) The full proof of all theorems; (ii) Additional details about the experiments and the experimental analysis; (iii) Extended related work section and discussions.

Marcolino, Jiang, and Tambe (2013) [14] study teams of agents that vote in discrete solution spaces. They show that a diverse team of weaker agents can overcome a uniform team made of copies of the best agent. However, this does not always occur, and they do not present ways to know when we should use diverse teams. Moreover, they lack a formal study of how the performance of diverse teams change as the number of agents and/or actions increases.

In this paper we shed new light on this problem, by presenting a new, more general model of diversity for teams of voting agents. Our model captures, better than the previous ones, the notion of a diverse team as a team of agents that tend to not agree on the same actions, and allows us to make new predictions. Our main insight is based on the notion of *spreading tail (ST)* and *non-spreading tail (NST)* agents. As we will show, a team of *ST* agents has a diverse behavior, i.e., they tend to not agree on the same actions. Hence, we can model a diverse team as a team of *ST* agents, and show that the performance improves as the size of the action space gets larger. We also prove upper and lower bounds on how fast different teams converge. The improvement can be large enough to overcome a uniform team of *NST* agents, even if individually the *ST* agents are weaker. As it is generally hard to find good solutions for problems with a large number of actions, it is important to know which teams to use in order to tackle such problems. Moreover, we show that the performance of a diverse team converges to the optimal one exponentially fast as the team grows. Our synthetic experiments provide even further insights about our model: even though the diverse team overcomes the uniform team in a large action space, the uniform team eventually will again play better than the diverse team as the action space keeps increasing if the best agent does not behave exactly like an *NST* agent.

Finally, we test our predictions by studying a system of voting agents, in the Go domain. We show that a uniform team made of copies of the best agent plays better in smaller boards, but is overcome by a diverse team as the board gets larger. We analyze the agents and verify that weak agents have a behavior closer to *ST* agents, while the best agent is closer to an *NST* agent. Therefore, we show that our predictions are verified in a real system, and can effectively be used while forming a multi-agent team.

2 Related Work

This paper is mainly related to team formation, but we also find related work in social choice and machine learning. We start by focusing on team formation research. Such study goes beyond computer science, and several works can be found in the economics literature. Hong and Page (2004) [8] is an impactful work showing the importance of diversity when forming (human) teams. Even though recently some of the mathematical arguments were put into question [19], it remains as a mile-stone on the study of the importance of diversity, as many researchers were influenced by their work. For example, LiCalzi and Surucu (2012) [11] present another model, that focuses on the importance of diversity when teams solve problems in large action spaces. However, both Hong and Page (2004) [8] and LiCalzi and Surucu (2012) [11] assume that the agents are able to know the utility of the solutions, and hence the team can pick the best solution found by one of its members. Therefore, their models do not apply for a team of vot-

ing agents. Lamberson and Page (2012) [10] study diversity in the context of forecasts. They assume that solutions are represented by real numbers, and a team converges to the average of the opinion of its members. Hence, they do not capture domains with discrete solutions, and the model also does not cover teams of voting agents.

In the multi-agent system literature, team formation is classically seen as selecting the team with maximum expected utility for a task, based on a model of the capabilities of each agent [16, 7]. However, in many domains we do not have such a model. The study of “ad-hoc” teamwork deals with multi-agent teams with absence of information [2, 1]. They focus, however, on how a new agent must decide its behavior in order to cooperate with agents of unknown types, not on picking the best team.

Recently, Marcolino, Jiang, and Tambe (2013) [14] showed the importance of diversity when forming teams of voting agents. They show that it is possible for a diverse team of weaker agents to overcome a uniform team of copies of the best agent, if the weaker agents are able to play better than the best agent at some world states. This is only a necessary condition, however, so it still does not provide ways to know when diverse or uniform teams should be used. Jiang *et al.* [9] propose a novel model to study diverse teams, where the agents’ votes are modeled as two samples from distributions: one that fixes the algorithm (or the biases) of the agent, and a second that models the actual voting process. Moreover, they experimentally study the performance of different voting rules in the Computer Go domain. However, the effects of changing the action space size are not studied, neither theoretically nor experimentally.

Concerning social choice, this paper is related to the view of voting as a way to discover an optimal choice (or ranking). Classical models study this view of voting for teams of identical agents [13, 5]. However, more recent works are also considering agents with different probability distribution functions. Caragiannis, Procaccia, and Shah (2013) [3] study which voting rules converge to the true ranking as the number of agents goes to infinity. In Soufiani, Parkes and Xia (2012) [18] the problem of inferring the true ranking is studied, assuming agents with different pdfs, but drawn from the same family. However, even though recent works on social choice are not assuming identical agents, they still do not provide a way to find the best teams of voting agents.

More related works can be found in machine learning. Ensemble systems are very common in machine learning, where a strong classifier is built by combining multiple weak classifiers, for example by voting [17]. Diversity is known to be important when forming an ensemble, and some systems try to minimize the correlation between the classifiers [4]. Still, an important problem is how to form the ensemble system, i.e., how to pick the classifiers that lead to the best predictions [6]. Our model, based on the notion of *spreading tail* and *non-spreading tail* agents, allows us to make predictions about teams as the action space and/or number of agents changes, and also compare the rate of change of the performance of different teams. To the best of our knowledge, there is no model similar to ours in the machine learning literature.

3 Model for Analysis of Diversity in Teams

Consider a problem defined by choosing an action a from a set of possible actions \mathbf{A} . Each a has an utility $U(a)$, and our goal is to maximize the utility. We always list the

actions in order from best to worst, therefore $U(a_j) > U(a_{j+1}) \forall j$ (a_0 is the best action). In some tasks (like in Section 4), a series of actions are chosen across different states, but here we focus on the decision process in a given state.

Consider a set of agents, voting to decide over actions. The agents do not know the utility of the actions, and vote for the action they believe to be the best according to their own decision procedure, characterized by a probability distribution (pdf). We write as $p_{i,j}$ the probability of agent i voting for action a_j . We denote by $p_{i,j}(m)$, when we explicitly talk about $p_{i,j}$ for an action space of size m . If the pdf of one agent is identical to the pdf of another agent, they will be referred to as copies of the same agent. The action that wins by plurality voting is taken by the team. Ties are broken randomly, except when we explicitly talk about a tie breaking rule. Let \mathbf{D}_m be the set of suboptimal actions ($a_j, j \neq 0$) assigned with a nonzero probability in the pdf of an agent i , and $d_m = |\mathbf{D}_m|$. We assume that there is a bound in the ratio of the suboptimal action with highest probability and the one with lowest nonzero probability, i.e., let $p_{i,min} = \min_{j \in \mathbf{D}_m} p_{i,j}$ and $p_{i,max} = \max_{j \in \mathbf{D}_m} p_{i,j}$; there is a constant α such that $p_{i,max} \leq \alpha p_{i,min} \forall$ agents i .

We define strength as the expected utility of an agent and/or a team. The probability of a team playing the best action will be called p_{best} . We first consider a setting where $U(a_0) \gg U(a_j) \forall j \neq 0$, hence we can use p_{best} as our measure of performance. We will later consider more general settings, where the first r actions have a high utility.

We define team formation as selecting from the space of all agents a limited number of agents that has the maximum strength by voting together to decide on actions. We study the effect of increasing the size m of the set of possible actions on the team formation problem. Intuitively, the change in team performance as m increases will be affected by how the pdf of the individual agents i change when m gets higher. As we increase m , d_m can increase or not change. Hence, we classify the agents as *spreading tail (ST)* agents or *non-spreading tail agents (NST)*.

We define *ST* agents as agents whose d_m is non-decreasing on m and $d_m \rightarrow \infty$ as $m \rightarrow \infty$. We consider that there is a constant $\epsilon > 0$, such that for all *ST* agents i , $\forall m$, $p_{i,0} \geq \epsilon$. We assume that $p_{i,0}$ does not change with m , although later we discuss what happens when $p_{i,0}$ changes.

We define *NST* agents as agents whose pdf does not change as the number of actions m increases. Hence, let m_{i0} be the minimum number of actions necessary to define the pdf of an *NST* agent i . We have that $\forall m, m' \geq m_{i0}, \forall j \leq m_{i0} p_{i,j}(m) = p_{i,j}(m')$, $\forall j > m_{i0} p_{i,j}(m) = 0$.

We first give an intuitive description of the concept of diversity, then define formally diverse teams. By diversity, we mean agents that tend to disagree. In Marcolino, Jiang, and Tambe (2013) [14], a diverse team is defined as a set of agents with different pdfs. Hence, they disagree because of having different probabilities of playing certain actions. Here, we generalize their definition to capture cases where agents disagree on actions, regardless of whether their pdfs are the same or not. Formally, we define a diverse team to be one consisting of a set of *ST* agents (either different *ST* agents or copies of the same *ST* agent). In our theoretical development we will show that this definition captures the notion of diversity: a team of *ST* agents will tend to not agree on the same suboptimal actions. We call uniform team as the team composed by copies of an *NST*

agent. This is an idealization to perform our initial analysis. We will later discuss more complex domains, where the agents of the uniform team also behave like *ST* agents.

We start with an example, to give an intuition about our model. Consider the agents in Table 1(a), where we show the pdf of the agents, and p_{best} of the uniform team (three copies of agent 1) and the diverse team (one copy of each agent). We assume agent 1 is an *NST* agent, while agent 2 and 3 are *ST* agents. In this situation the uniform team plays better than the diverse team. Now let's add one more action to the problem. Because agent 2 and 3 are *ST* agents, the probability mass on action 2 scatters to the newly added action (Table 1(b)). Hence, while before the *ST* agents would always agree on the same suboptimal action if they both did not vote for the optimal action, now they might vote for different suboptimal actions, creating a tie between each suboptimal action and the optimal one. Because ties are broken randomly, when this happens there will be a $1/3$ chance that the tie will be broken in favor of the optimal action. Hence, p_{best} increases when the probability of the *ST* agents agreeing on the same suboptimal actions decreases, and the diverse team now plays better than the uniform team, even though individually agents 2 and 3 are weaker than agent 1.

Table 1: Performance of diverse team increases when the number of actions increases.

(a) With 2 actions, uniform team plays better than diverse team.

| Agents | Action 1 | Action 2 |
|----------------------|----------|----------|
| Agent 1 | 0.6 | 0.4 |
| Agent 2 | 0.55 | 0.45 |
| Agent 3 | 0.55 | 0.45 |
| Uniform p_{best} : | 0.648 | |
| Diverse p_{best} : | 0.599 | |

(b) When we add one more action, diverse team plays better than uniform team.

| Agents | Action 1 | Action 2 | Action 3 |
|----------------------|----------|----------|----------|
| Agent 1 | 0.6 | 0.4 | 0 |
| Agent 2 | 0.55 | 0.25 | 0.2 |
| Agent 3 | 0.55 | 0.15 | 0.3 |
| Uniform p_{best} : | 0.648 | | |
| Diverse p_{best} : | 0.657 | | |

We now present our theoretical work. First we show that the performance of a diverse team converges when $m \rightarrow \infty$, to a value that is higher than the performance for any other m .

Theorem 1. $p_{best}(m)$ of a diverse team of n agents converges to a certain value \tilde{p}_{best} as $m \rightarrow \infty$. Furthermore, $\tilde{p}_{best} \geq p_{best}(m)$, $\forall m$.

Proof. Let $p_{i,min} = \min_{j \in \mathbf{D}_m} p_{i,j}$, $p_{i,max} = \max_{j \in \mathbf{D}_m} p_{i,j}$ and \mathbf{T} be the set of agents in the team. By our assumptions, there is a constant α such that $p_{i,max} \leq \alpha p_{i,min}$ for all agents i . Then, we have that $1 \geq 1 - p_{i,0} = \sum_{j \in \mathbf{D}_m} p_{i,j} \geq d_m p_{i,min}$. Therefore, $p_{i,min} \leq \frac{1}{d_m} \rightarrow 0$ as d_m tends to ∞ with m . Similarly, $\alpha p_{i,min} \rightarrow 0$ as $d_m \rightarrow \infty$. As $p_{i,j} \leq \alpha p_{i,min}$ we have that $\forall j p_{i,j} \rightarrow 0$ as $d_m \rightarrow \infty$. We show that this implies that when $m \rightarrow \infty$, weak agents never agree on the same suboptimal action. Let i_1 and i_2 be two arbitrary agents. Without loss of generality, assume i_2 's d_m ($d_m^{(i_2)}$) is greater than or equal i_1 's d_m ($d_m^{(i_1)}$). The probability (σ_{i_1,i_2}) of i_1 and i_2 agreeing on the same suboptimal action is upper bounded by $\sigma_{i_1,i_2} = \sum_{a_j \in \mathbf{A} \setminus a_0} p_{i_1,j} p_{i_2,j} \leq d_m^{(i_2)} p_{i_1,max} p_{i_2,max} \leq$

$d_m^{(i_2)} \alpha p_{i_2, \min} p_{i_1, \max} \leq \alpha p_{i_1, \max}$ (as $d_m^{(i_2)} p_{i_2, \min} \leq 1$). We have that $\alpha p_{i_1, \max} \rightarrow 0$ as $p_{i_1, \max} \rightarrow 0$, because α is a constant. Hence the probability of any two agents agreeing on a suboptimal action is $\frac{\sum_{i_1 \in \mathbb{T}} \sum_{i_2 \in \mathbb{T}, i_2 \neq i_1} \sigma_{i_1, i_2}}{2} \leq \frac{n(n-1)}{2} \max_{i_1, i_2} \sigma_{i_1, i_2} \rightarrow 0$, as n is a constant.

Hence, when $m \rightarrow \infty$, the diverse team only chooses a suboptimal action if all agents vote for a different suboptimal action or in a tie between the optimal action and suboptimal actions (because ties are broken randomly). Therefore, p_{best} converges to:

$$\tilde{p}_{best} = 1 - \prod_{i=1}^n (1 - p_{i,0}) - \sum_{i=1}^n (p_{i,0} \prod_{j=1, j \neq i}^n (1 - p_{j,0})) \frac{n-1}{n}, \quad (1)$$

that is, the total probability minus the cases where the best action is not chosen: the second term covers the case where all agents vote for a suboptimal action and the third term covers the case where one agent votes for the optimal action and all other agents vote for suboptimal actions.

When m is finite, the agents might choose a suboptimal action by agreeing over that suboptimal action. Therefore, we have that $p_{best}(m) \leq \tilde{p}_{best} \forall m$. \square

Let $p_{best}^{uniform}(m)$ be p_{best} of the uniform team, with m actions. A uniform team is not affected by increasing m , as the pdf of an *NST* agent will not change. Hence, $p_{best}^{uniform}(m)$ is the same, $\forall m$. If \tilde{p}_{best} is high enough so that $\tilde{p}_{best} \geq p_{best}^{uniform}(m)$, the diverse team will overcome the uniform team, when $m \rightarrow \infty$. Therefore, the diverse team will be better than the uniform team when m is large enough.

In practice, a uniform team made of copies of the best agent might not behave exactly like a team of *NST* agents, as the best agent could also increase its d_m as m gets larger. We discuss this situation in Section 4. In order to perform that study, we derive in the following corollary how fast p_{best} converges to \tilde{p}_{best} , as a function of d_m .

Corollary 1. $p_{best}(m)$ of a diverse team increases to \tilde{p}_{best} in the order of $O\left(\frac{1}{d_m^{min}}\right)$ and $\Omega\left(\frac{1}{d_m^{max}}\right)$, where d_m^{max} is the highest and d_m^{min} the lowest d_m of the team.

Proof. We assume here the notation that was used in the previous proof. First we show a lowerbound on $p_{best}(m)$. We have that $p_{best}(m) = 1 - \psi_1$, where ψ_1 is the probability of the team picking a suboptimal action. $\psi_1 = \psi_2 + \psi_3$, where ψ_2 is the probability of no agent agreeing and the team picks a suboptimal action and ψ_3 is the probability of at least two agents agreeing and the team picks a suboptimal action. Hence, $p_{best}(m) = 1 - \psi_2 - \psi_3 = \tilde{p}_{best} - \psi_3 \geq \tilde{p}_{best} - \psi_4$, where ψ_4 is the probability of at least two agents agreeing. Let $\sigma^{max} = \max_{i_1, i_2} \sigma_{i_1, i_2}$, and i_1^* and i_2^* are the agents whose $\sigma_{i_1^*, i_2^*} = \sigma^{max}$. We have that:

$$\begin{aligned} p_{best}(m) &\geq \tilde{p}_{best} - \frac{n(n-1)}{2} \sigma^{max} \geq \tilde{p}_{best} - \frac{n(n-1)}{2} d_m^{(i_2^*)} p_{i_1^*, \max} p_{i_2^*, \max} \\ &\geq \tilde{p}_{best} - \frac{n(n-1)}{2} d_m^{(i_2^*)} \alpha p_{i_1^*, \min} \alpha p_{i_2^*, \min} \geq \tilde{p}_{best} - \frac{n(n-1)}{2} \alpha^2 \frac{1}{d_m^{(i_1^*)}}, \end{aligned}$$

where the last inequality holds since $p_{i, \min} \leq \frac{1}{d_m}$. Hence, $p_{best}(m) \geq \tilde{p}_{best} - \frac{n(n-1)}{2} \alpha^2 \frac{1}{d_m^{min}} \rightsquigarrow \tilde{p}_{best} - p_{best}(m) \leq O\left(\frac{1}{d_m^{min}}\right)$.

Now we show an upper bound: $p_{best}(m) = \tilde{p}_{best} - \psi_3 \leq \tilde{p}_{best} - \psi_5$, where ψ_5 is the probability of at least two agents agreeing and no agents vote for the optimal action. Let $\sigma^{min} = \min_{i_1, i_2} \sigma_{i_1, i_2}$; i_1^* and i_2^* are the agents whose $\sigma_{i_1^*, i_2^*} = \sigma^{min}$; and $p_{max,0} = \max_{i \in \mathbf{T}} p_{i,0}$. Without loss of generality, we assume that $d_m^{(i_2^*)} \geq d_m^{(i_1^*)}$. Therefore:

$$\begin{aligned}
 p_{best}(m) &\leq \tilde{p}_{best} - \frac{n(n-1)}{2} \sigma^{min} (1 - p_{max,0})^{n-2} \\
 &\leq \tilde{p}_{best} - \frac{n(n-1)}{2} d_m^{(i_1^*)} p_{i_1^*, min} p_{i_2^*, min} (1 - p_{max,0})^{n-2} \\
 &\leq \tilde{p}_{best} - \frac{n(n-1)}{2} d_m^{(i_1^*)} \frac{p_{i_1^*, max} p_{i_2^*, max}}{\alpha^2} (1 - p_{max,0})^{n-2} \\
 &\leq \tilde{p}_{best} - \frac{n(n-1)}{2} \alpha^{-2} \frac{1}{d_m^{i_2^*}} (1 - p_{max,0})^{n-2} \\
 &\leq \tilde{p}_{best} - \frac{n(n-1)}{2} \alpha^{-2} \frac{1}{d_m^{max}} (1 - p_{max,0})^{n-2} \\
 &\rightsquigarrow \tilde{p}_{best} - p_{best}(m) \geq \Omega\left(\frac{1}{d_m^{max}}\right). \square
 \end{aligned}$$

Hence, agents that change their d_m faster will converge faster to \tilde{p}_{best} . This is an important result when we consider later more complex scenarios where the d_m of the agents of the uniform team also change.

Note that \tilde{p}_{best} depends on the number of agents n (Equation 1). Now we show that the diverse team tends to always play the optimal action, as $n \rightarrow \infty$.

Theorem 2. *\tilde{p}_{best} converges to 1, as $n \rightarrow \infty$. Furthermore, $1 - \tilde{p}_{best}$ converges exponentially to 0, that is, \exists constant c , such that $1 - \tilde{p}_{best} \leq c(1 - \frac{\epsilon}{2})^n$, $\forall n \geq \frac{2}{\epsilon}$. However, the performance of the uniform team improves as $n \rightarrow \infty$ only if $p_{s,0} = \max_j p_{s,j}$, where s is the best agent.*

Proof. By the previous proof, we know that when $m \rightarrow \infty$ the diverse team plays the optimal action with probability given by \tilde{p}_{best} . We show that $1 - \tilde{p}_{best} \rightarrow 0$ exponentially as $n \rightarrow \infty$ (this naturally induces $\tilde{p}_{best} \rightarrow 1$). We first compute an upper bound for $\sum_{i=1}^n p_{i,0} \prod_{j=1, j \neq i}^n (1 - p_{j,0})$:

$$\begin{aligned}
 \sum_{i=1}^n p_{i,0} \prod_{j=1, j \neq i}^n (1 - p_{j,0}) &\leq \sum_{i=1}^n p_{i,0} (1 - p_{min,0})^{n-1} \leq n p_{max,0} (1 - p_{min,0})^{n-1} \\
 &\leq n(1 - \epsilon)^{n-1} \text{ for } p_{max,0} = \max_i p_{i,0}, p_{min,0} = \min_j p_{j,0}
 \end{aligned}$$

Since $\prod_{i=1}^n (1 - p_{i,0}) \leq (1 - \epsilon)^n$, thus we have that $1 - \tilde{p}_{best} \leq (1 - \epsilon)^n + n(1 - \epsilon)^{n-1}$. So we only need to prove that there exists a constant c such that $(1 - \epsilon)^n + n(1 - \epsilon)^{n-1} \leq c(1 - \frac{\epsilon}{2})^n$, as follows:

$$\begin{aligned}
 \frac{(1 - \epsilon)^{n+1} + (n+1)(1 - \epsilon)^n}{(1 - \epsilon)^n + n(1 - \epsilon)^{n-1}} &= (1 - \epsilon) \frac{1 - \epsilon + n + 1}{1 - \epsilon + n} = 1 - \epsilon + \frac{1 - \epsilon}{1 - \epsilon + n} \\
 &\leq 1 - \frac{1}{2}\epsilon, \text{ if } n \geq \frac{2}{\epsilon} \text{ (by setting } \frac{1 - \epsilon}{1 - \epsilon + n} \leq \frac{\epsilon}{2} \text{)}.
 \end{aligned}$$

Hence, $\exists c$, such that $(1 - \epsilon)^n + n(1 - \epsilon)^{n-1} \leq c(1 - \frac{\epsilon}{2})^n$ when $n \geq \frac{2}{\epsilon}$. Therefore, the performance converges exponentially.

For the uniform team, the probability of playing the action that has the highest probability in the pdf of the best agent converges to 1 as $n \rightarrow \infty$ [13]. Therefore, the performance only increases as $n \rightarrow \infty$ if the optimal action is the one that has the highest probability. \square

Now we show that we can achieve further improvement in a diverse team by breaking ties in favor of the strongest agent.

Theorem 3. *When $m \rightarrow \infty$, breaking ties in favor of the strongest agent is the optimal tie-breaking rule for a diverse team.*

Proof. Let s be one of the agents. If we break ties in favor of s , the probability of voting for the optimal choice will be given by:

$$\tilde{p}_{best} = 1 - \prod_{i=1}^n (1 - p_{i,0}) - (1 - p_{s,0}) \left(\sum_{i=1, i \neq s}^n p_{i,0} \prod_{j=1, j \neq i, j \neq s}^n (1 - p_{j,0}) \right) \quad (2)$$

It is clear that Equation 2 is maximized by choosing agent s with the highest $p_{s,0}$. However, we still have to show that it is better to break ties in favor of the strongest agent than breaking ties randomly. That is, we have to show that Equation 2 is always higher than Equation 1.

Equation 2 differs from Equation 1 only on the last term. Therefore, we have to show that the last term of Equation 2 is smaller than the last term of Equation 1. Let's begin by rewriting the last term of Equation 1 as:

$$\begin{aligned} & \frac{n-1}{n} \sum_{i=1}^n p_{i,0} \prod_{j=1, j \neq i}^n (1 - p_{j,0}) = \\ & \frac{n-1}{n} (1 - p_{s,0}) \sum_{i=1, i \neq s}^n p_{i,0} \prod_{j=1, j \neq i, j \neq s}^n (1 - p_{j,0}) + \frac{n-1}{n} p_{s,0} \prod_{j=1, j \neq s}^n (1 - p_{j,0}) \end{aligned}$$

This implies that:

$$\frac{n-1}{n} \sum_{i=1}^n p_{i,0} \prod_{j=1, j \neq i}^n (1 - p_{j,0}) \geq \frac{n-1}{n} (1 - p_{s,0}) \sum_{i=1, i \neq s}^n p_{i,0} \prod_{j=1, j \neq i, j \neq s}^n (1 - p_{j,0}).$$

We know that:

$$\begin{aligned} & (1 - p_{s,0}) \sum_{i=1, i \neq s}^n p_{i,0} \prod_{j=1, j \neq i, j \neq s}^n (1 - p_{j,0}) = \\ & \frac{n-1}{n} (1 - p_{s,0}) \sum_{i=1, i \neq s}^n p_{i,0} \prod_{j=1, j \neq i, j \neq s}^n (1 - p_{j,0}) + \\ & \frac{1}{n} (1 - p_{s,0}) \sum_{i=1, i \neq s}^n p_{i,0} \prod_{j=1, j \neq i, j \neq s}^n (1 - p_{j,0}) \end{aligned}$$

Therefore, for the last term of Equation 2 to be smaller than the last term of Equation 1 we have to show that:

$$\frac{n-1}{n} p_{s,0} \prod_{j=1, j \neq s}^n (1 - p_{j,0}) \geq \frac{1}{n} (1 - p_{s,0}) \sum_{i=1, i \neq s}^n p_{i,0} \prod_{j=1, j \neq s, j \neq i}^n (1 - p_{j,0})$$

It follows that the previous equation will be true if:

$$\begin{aligned} p_{s,0} & \geq (1 - p_{s,0}) \frac{\sum_{i=1, i \neq s}^n p_{i,0} \prod_{j=1, j \neq i, j \neq s}^n (1 - p_{j,0})}{(n-1) \prod_{j=1, j \neq s}^n (1 - p_{j,0})} \\ p_{s,0} & \geq (1 - p_{s,0}) \frac{1}{n-1} \sum_{i=1, i \neq s}^n \frac{p_{i,0}}{(1 - p_{i,0})} \\ \frac{p_{s,0}}{(1 - p_{s,0})} & \geq \frac{\sum_{i=1, i \neq s}^n \frac{p_{i,0}}{(1 - p_{i,0})}}{n-1} \end{aligned}$$

As s is the strongest agent the previous inequality is always true. This is because $\frac{p_{s,0}}{1-p_{s,0}} = \frac{\sum_{i=1, i \neq s}^n \frac{p_{s,0}}{(1-p_{s,0})}}{n-1}$ and $\frac{p_{s,0}}{1-p_{s,0}} \geq \frac{p_{i,0}}{(1-p_{i,0})} \forall i \neq s$. Therefore, it is always better to break ties in favor of the strongest agent than breaking ties randomly. \square

Next we show that with one additional assumption, not only the diverse team converges to \tilde{p}_{best} , but also p_{best} monotonically increases with m . Our additional assumption is that higher utility actions have higher probabilities, i.e., if $U(a_j) \geq U(a_{j'})$, then $p_{i,j} \geq p_{i,j'}$.

Theorem 4. *The performance of a diverse team monotonically increases with m , if $U(a_j) \geq U(a_{j'})$ implies that $p_{i,j} \geq p_{i,j'}$.*

Proof. Let an event be one voting iteration, where each agent from a set votes for an action. We denote by $P(\mathbf{V})$ the probability of occurrence of any event in \mathbf{V} (hence, $P(\mathbf{V}) = \sum_{v \in \mathbf{V}} p(v)$). We call it a winning event if in the event the action chosen by plurality is the best action a_0 (including ties). We assume that for all agents i , if $U(a_j) \geq U(a_{j'})$, then $p_{i,j} \geq p_{i,j'}$.

We show by mathematical induction that we can divide the probability of multiple suboptimal actions into a new action and $p_{best}(m+1) \geq p_{best}(m)$. Let λ be the number of actions whose probability is being divided. The base case holds trivially when $\lambda = 0$. That is, there is a new action, but all agents have a 0 probability of voting for that new action. In this case we have that p_{best} does not change, therefore $p_{best}(m+1) \geq p_{best}(m)$.

Now assume that we divided the probability of λ actions and it is true that $p_{best}(m+1) \geq p_{best}(m)$. We show that it is also true for $\lambda + 1$. Hence, let's pick one more action to divide the probability. Without loss of generality, assume it is action a_{d_m} , for agent c , and its probability is being divided into action a_{d_m+1} . Therefore, $p'_{c,d_m} = p_{c,d_m} - \beta$ and $p'_{c,d_m+1} = p_{c,d_m+1} + \beta$, for $0 \leq \beta \leq p_{c,d_m}$. Let $p_{best}^{after}(m+1)$ be the probability of voting for the best action after this new division, and $p_{best}^{before}(m+1)$ the probability before this new division. We show that $p_{best}^{after}(m+1) \geq p_{best}^{before}(m+1)$.

Let Γ be the set of all events where all agents voted, except for agent c (the order does not matter, so we can consider agent c is the last one to post its vote). If $\gamma \in \Gamma$ will be a winning event no matter if agent c votes for a_{d_m} or a_{d_m+1} , then changing agent c 's pdf will not affect the probability of these winning events. Hence, let $\Gamma' \subset \Gamma$ be the set of all events that will become a winning event depending if agent c does not vote for a_{d_m} or a_{d_m+1} . Given that $\gamma \in \Gamma'$ already happened, the probability of winning or losing is equal to the probability of agent c not voting for a_{d_m} or a_{d_m+1} .

Now let's divide Γ' in two exclusive subsets: $\Gamma_{d_m+1} \subset \Gamma'$, where for each $\gamma \in \Gamma_{d_m+1}$ action a_{d_m+1} is in tie with action a_0 , so if agent c does not vote for a_{d_m+1} , γ will be a winning event; $\Gamma_{d_m} \subset \Gamma'$, where for each $\gamma \in \Gamma_{d_m}$ action a_{d_m} is in tie with action a_0 , so if agent c does not vote for a_{d_m} , γ will be a winning event. We do not consider events where both a_{d_m+1} and a_{d_m} are in tie with a_0 , as in that case the probability of a winning event does not change (it is given by $1 - p'_{c,d_m} - p'_{c,d_m+1} = 1 - p_{c,d_m} - p_{c,d_m+1}$).

Note that for each $\gamma \in \Gamma_{d_m+1}$, the probability of a winning event equals $1 - p'_{c,d_m+1}$. Therefore, after changing the pdf of agent c , for each $\gamma \in \Gamma_{d_m+1}$, the probability of a winning event decreases by β . Similarly, for each $\gamma \in \Gamma_{d_m}$, the probability of a winning event equals $1 - p'_{c,d_m}$. Therefore, after changing the pdf of agent c , for each $\gamma \in \Gamma_{d_m}$, the probability of a winning event increases by β .

Therefore, $p_{best}^{after}(m+1) \geq p_{best}^{before}(m+1)$ if and only if $P(\Gamma_{d_m}) \geq P(\Gamma_{d_m+1})$. Note that $\forall \gamma \in \Gamma_{d_m+1}$ there are more agents that voted for a_{d_m+1} than for a_{d_m} . Also, $\forall \gamma \in \Gamma_{d_m}$ there are more agents that voted for a_{d_m} than for a_{d_m+1} . If, for all agents i , $p_{i,d_m} \geq p_{i,d_m+1}$, we have that $P(\Gamma_{d_m}) \geq P(\Gamma_{d_m+1})$. Therefore, $p_{best}^{after}(m+1) \geq p_{best}^{before}(m+1)$, so we still have that $p_{best}(m+1) \geq p_{best}(m)$. Also note that in order for the next step of the induction to still be valid, so that we can still divide the probability of one more action, it is necessary that $p'_{c,d_m} \geq p'_{c,d_m+1}$. \square

In the previous theorems we focused on the probability of playing the best action, assuming that $U(a_0) \gg U(a_j) \forall j \neq 0$. We show now that the theorems still hold in more general domains where r actions ($\mathbf{A}_r \subset \mathbf{A}$) have a significant high utility, i.e., $U(a_{j_1}) \gg U(a_{j_2}) \forall j_1 < r, j_2 \geq r$. Hence, we now focus on the probability of playing any action in \mathbf{A}_r . We assume that our assumptions are also generalized, i.e., $p_{i,j} > \epsilon \forall j < r$, and the number d_m of suboptimal actions ($a_j, j \geq r$) in the \mathbf{D}_m set increases with m for ST agents.

Theorem 5. *The previous theorems generalize to settings where $U(a_{j_1}) \gg U(a_{j_2}) \forall j_1 < r, j_2 \geq r$.*

Proof Sketch We give here a proof sketch. We just have to generate new pdfs $p'_{i,j}$, such that $p'_{i0} = \sum_{j=0}^{r-1} p_{i,j}$, and $p'_{i,b} = p_{i,b+r-1}, \forall b \neq 0$. We can then reapply the proofs of the previous theorems, but replacing $p_{i,j}$ by $p'_{i,j}$. Note that this does not guarantee that all agents will tend to agree on the same action in \mathbf{A}_r ; but the team will still tend to pick any action in \mathbf{A}_r , since the agents are more likely to agree on actions in \mathbf{A}_r than on actions in $\mathbf{A} \setminus \mathbf{A}_r$. \square

Now we discuss a different generalization: what happens when $p_{i,0}$ decreases as m increases (\forall agents i). If $p_{i,0} \rightarrow \tilde{p}_{i,0}$ as $m \rightarrow \infty$, the performance in the limit for a diverse team will be \tilde{p}_{best} evaluated at $\tilde{p}_{i,0}$. Moreover, even if $p_{i,0} \rightarrow 0$, our conclusions about relative team performance are not affected as long as we are comparing two ST teams that have similar $p_{i,0}$: the same argument as in Corollary 1 implies that the team with faster growing d_m will perform better.

4 Experimental Analysis

4.1 Synthetic Experiments

We present synthetic experiments, in order to better understand what happens in real systems. We generate agents by randomly creating pdfs and calculate the probability of playing the best action (p_{best}) of the generated teams. We use a uniform distribution to generate all random numbers. When creating a pdf, we rescale the values assigned randomly, so that the overall sum of the pdf is equal to 1.

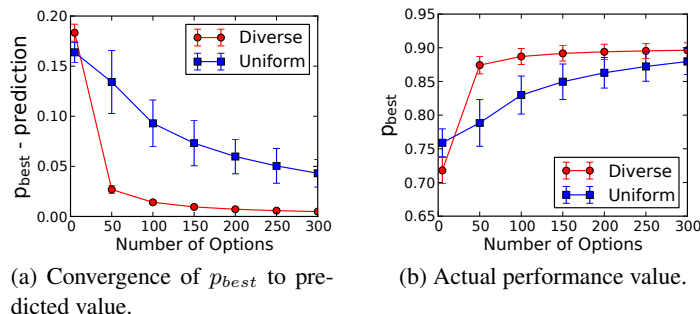


Fig. 1: Comparing diverse and uniform when uniform also increases d_m .

As we said earlier, uniform teams composed by *NST* agents is an idealization. In more complex domains, the best agent will not behave exactly like an *NST* agent; the number of suboptimal actions with a non-zero probability (d_m) will also increase as the action space gets larger. We perform synthetic experiments to study this situation. We consider that the best agent is still closer to an *NST* agent, therefore it increases its d_m at a slower rate than the agents of the diverse team.

In our first experiment, we use teams of 4 agents. For each agent of the diverse team, $p_{i,0}$ is chosen randomly between 0.6 and 0.7. The remaining is distributed randomly from 10% to 20% of the next best actions (the number of actions that will receive a positive probability is also decided randomly). For the uniform team, we make copies of the best agent (with highest $p_{i,0}$) of the diverse team, but distribute the remaining probability randomly from 1% to 3% of the next best actions.

We can see the average result for 200 random teams in Figure 1, where in Figure 1(a) we show the difference between the performance in the limit (\tilde{p}_{best}) and the actual $p_{best}(m)$ for the diverse and the uniform teams; in Figure 1(b) we show the average $p_{best}(m)$ of the teams. As can be seen, when the best agents increase their d_m at a slower rate than the agents of the diverse team, the uniform teams converge slower to \tilde{p}_{best} . Even though they play better than the diverse teams for a small m , they are surpassed by the diverse teams as m increases. However, because \tilde{p}_{best} of the uniform teams is actually higher than the one of the diverse teams, eventually the performance of the uniform teams get closer to the performance of the diverse teams, and will be better than the one of the diverse teams again for a large enough m .

This situation is expected according to Theorem 1. If the d_m of the best agent also increases as m gets larger, the uniform team will actually behave like a diverse team and also converge to \tilde{p}_{best} . $\tilde{p}_{best}^{uniform} \geq \tilde{p}_{best}^{diverse}$, as the best agent has a higher probability of playing the optimal action. Hence, in the limit the uniform team will play better than the diverse team. However, as we saw in Corollary 1, the speed of convergence is in the order of $1/d_m$. There-




Fig. 2: p_{best} of a diverse team as the number of agents increases.

fore, the diverse team will converge faster, and can overcome the uniform team for moderately large m .

As Theorem 2 only holds when $m \rightarrow \infty$, we also explore the effect of increasing the number of agents for a large m . The \tilde{p}_{best} of a team of agents is shown as the dashed line in Figure 2. We are plotting for agents that have a probability of playing the best action of only 0.1, but as we can see the probability quickly grows as the number of agents increases. We also calculate p_{best} for random teams from 2 to 6 agents (shown as the continuous line), when there are 300 available actions. Each agent has a probability of playing the best action of 0.1, and the remaining probability is randomly distributed over the 10% next best actions. As can be seen, the teams have a close performance to the expected. We only show up to 6 agents because it is too computationally expensive to calculate the pdfs of larger teams.

4.2 Computer Go

We present now results in a real system. We use in our experiments 4 different Go software: Fuego 1.1, GnuGo 3.8, Pachi 9.01, MoGo 4, and two (weaker) variants of Fuego (Fuego Δ and Fuego Θ), in a total of 6 different, publicly available, agents. Fuego is considered the strongest agent among all of them. Fuego is an implementation of the UCT Monte Carlo Go algorithm, therefore it uses heuristics to simulate games in order to evaluate board configurations. Fuego uses mainly 5 heuristics during these simulations, and they are executed in a hierarchical order. The original Fuego agent follows the order <Atari Capture, Atari Defend, Lowlib, Pattern> (the heuristic called Nakade is not enabled by default). Our variation called Fuego Δ follows the order <Atari Defend, Atari Capture, Pattern, Nakade, Lowlib>, while Fuego Θ follows the order <Atari Defend, Nakade, Pattern, Atari Capture, Lowlib>. Also, Fuego Δ and Fuego Θ have half of the memory available when compared with the original Fuego.

All our results are obtained by playing either 1000 games (to evaluate individual agents) or 2000 games (to evaluate teams), in a HP dl165 with dual dodeca core, 2.33GHz processors and 48GB of RAM. We compare results obtained by playing against a fixed opponent. Therefore, we evaluate systems playing as white, against the original Fuego playing as black. We removed all databases and specific board size knowledge of the agents, including the opponent. We call Diverse as the team composed of all 6 agents, and Uniform as the team composed of 6 copies of Fuego. Each agent is initialized with a different random seed, therefore they will not vote for the same action all the time in a given world state, due to the characteristics of the search algorithms. In all the graphs we present in this section, the error bars show the confidence interval, with 99% of confidence ($p = 0.01$).

We evaluate the performance of the teams over 7 different board sizes. We changed the time settings of individual agents as we increased the board size, in order to keep their strength as constant as possible. The average winning rates of the team members is shown in Table 2, while Table 3 show the winning rates of the individual agents.²

² In our first experiment, Diverse improved from 56.1% on 9x9 to 85.9% on 19x19. We noted, however, that some of the diverse agents were getting stronger in relation to the opponent as the board size increased. Hence, by changing the time setting to keep the strength constant, we are actually making our claims harder to show, not easier.

Table 2: Average winning rates of the team members across different board sizes. Note that these are not the winning rates of the teams.

| Team | 9x9 | 11x11 | 13x13 | 15x15 | 17x17 | 19x19 | 21x21 |
|---------|-------|-------|-------|-------|-------|-------|-------|
| Diverse | 32.2% | 30.8% | 29.6% | 29.4% | 31.5% | 31.9% | 30.3% |
| Uniform | 48.1% | 48.6% | 46.1% | 48.0% | 49.3% | 46.9% | 46.6% |

Table 3: Winning rates of each one of the agents across different board sizes.

| Agent | 9x9 | 11x11 | 13x13 | 15x15 | 17x17 | 19x19 | 21x21 |
|----------------|-------|-------|-------|-------|-------|-------|-------|
| Fuego | 48.1% | 48.6% | 46.1% | 48.0% | 49.3% | 46.9% | 46.6% |
| GnuGo | 1.1% | 1.1% | 1.9% | 1.9% | 4.5% | 6.8% | 6.1% |
| Pachi | 25.7% | 22.9% | 25.8% | 26.9% | 23.5% | 20.8% | 11.0% |
| MoGo | 27.6% | 26.4% | 22.7% | 22.0% | 27.1% | 30.1% | 27.1% |
| Fuego Δ | 45.7% | 45.8% | 42.2% | 40.4% | 43.0% | 44.5% | 47.4% |
| Fuego Θ | 45.5% | 40.2% | 39.2% | 37.6% | 41.8% | 42.3% | 43.6% |

We can see our results in Figure 4 (a). Diverse improves from 58.1% on 9x9 to 72.1% on 21x21, an increase in winning rate that is statistically significant with $p < 2.2 \times 10^{-16}$. This result is expected according to Theorem 1. Uniform changes from 61.0% to 65.8%, a statistically significant improvement with $p = 0.0018$. As we saw before, an increase in the performance of Uniform can also be expected, as the best agent might not be a perfect *NST* agent. A linear regression of the results of both teams gives a slope of 0.010 for the diverse team (adjusted R^2 : 0.808, $p = 0.0036$) and 0.005 for the uniform team (adjusted R^2 : 0.5695, $p = 0.0305$). Therefore, the diverse team improves its winning rate faster than the uniform team. To check if this is a significant difference, we evaluate the interaction term in a linear regression with multiple variables. We find that the influence of board size is higher on Diverse than on Uniform with $p = 0.0797$ (estimated coefficient of “size of the board \times group type”: -10.321 , adjusted R^2 : 0.7437). Moreover, on the 9x9 board Diverse is worse than Uniform ($p = 0.0663$), while on the 21x21 board Diverse is better with high statistical significance ($p = 1.941 \times 10^{-5}$). We also analyze the performance of the teams subtracted by the average strength of their members (Figure 4 (b)), in order to calculate the increase in winning rate achieved by “teamwork” and compensate fluctuations on the winning rate of the agents as we change the board size. Again, the diverse team improves faster than the uniform team. A linear regression results in a slope of 0.0104 for Diverse (adjusted R^2 : 0.5549, $p = 0.0546$) and 0.0043 for Uniform (adjusted R^2 : 0.1283, $p = 0.258$).

We also evaluate the performance of teams of 4 agents (Diverse 4 and Uniform 4). For Diverse 4, we removed Fuego Δ and Fuego Θ from the Diverse team. As can be seen in Figure 3, the impact of adding more agents is higher for the diverse team in a larger board size (21x21). In the 9x9 board, the difference between




Fig. 3: Winning rates for 4 and 6 agents teams.

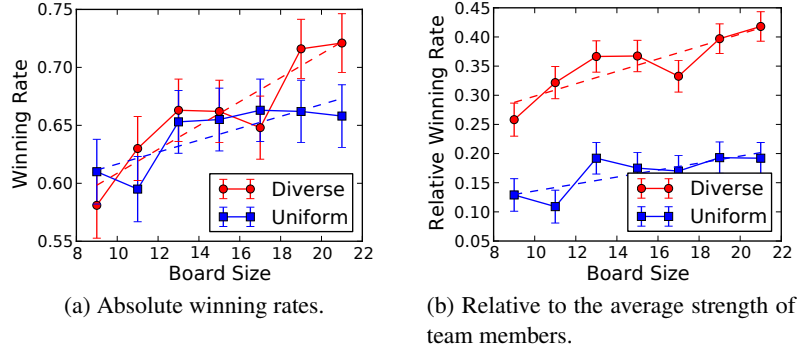


Fig. 4: Winning rate in the real Computer Go system.

Diverse 4 and Diverse 6 is only 4.4%; while in 21x21 it is 14%. Moreover, we can see a higher impact of adding agents for the diverse team, than for the uniform team. These results would be expected according to Theorem 2.

As can be seen, the predictions of our theory holds: the diverse team improves significantly as we increase the action space. The improvement is enough to make it change from playing worse than the uniform team on 9x9 to playing better than the uniform team with statistical significance on the 21x21 board. Furthermore, we show a higher impact of adding more agents when the size of the board is larger.

4.3 Analysis

To test the assumptions of our model, we estimate a pdf for each one of the agents. For each board size, and for each one of 1000 games from our experiments, we randomly choose a board state between the first and the last movement. We make Fuego evaluate the chosen board, but with a time limit 50x higher than the default one. Therefore, we use this much stronger version of Fuego to approximate the true ranking of all actions. For each board size, we run all agents in each board sample and check in which position of the approximated true ranking they play. This allow us to build a histogram for each agent and board size combination. Some examples can be seen in Figure 5. We can see that a strong agent, like Fuego, has most of its probability mass on the higher ranked actions, while weaker agents, like GnuGo, has the mass of its pdf distributed over a larger set of actions, creating a larger tail. Moreover, the probability mass of GnuGo is spread over a larger number of actions when we increase the size of the board.

We study how the pdfs of the agents change as we increase the action space. Our hypothesis is that weaker agents will have a behavior closer to *ST* agents, while stronger agents to *NST* agents. In Figure 6 (a) we show how many actions receive a probability higher than 0. As can be seen, Fuego does not behave exactly like an *NST* agent. However, it does have a slower growth rate than the other agents. A linear regression gives the following slopes: 13.08, 19.82, 19.05, 15.82, 15.69, 16.03 for Fuego, Gnugo, Pachi, Mogo, Fuego Δ and Fuego Θ , respectively (R^2 : 0.95, 0.98, 0.94, 0.98, 0.98, 0.98, respectively). It is clear, therefore, that the probability mass of weak agents is distributed

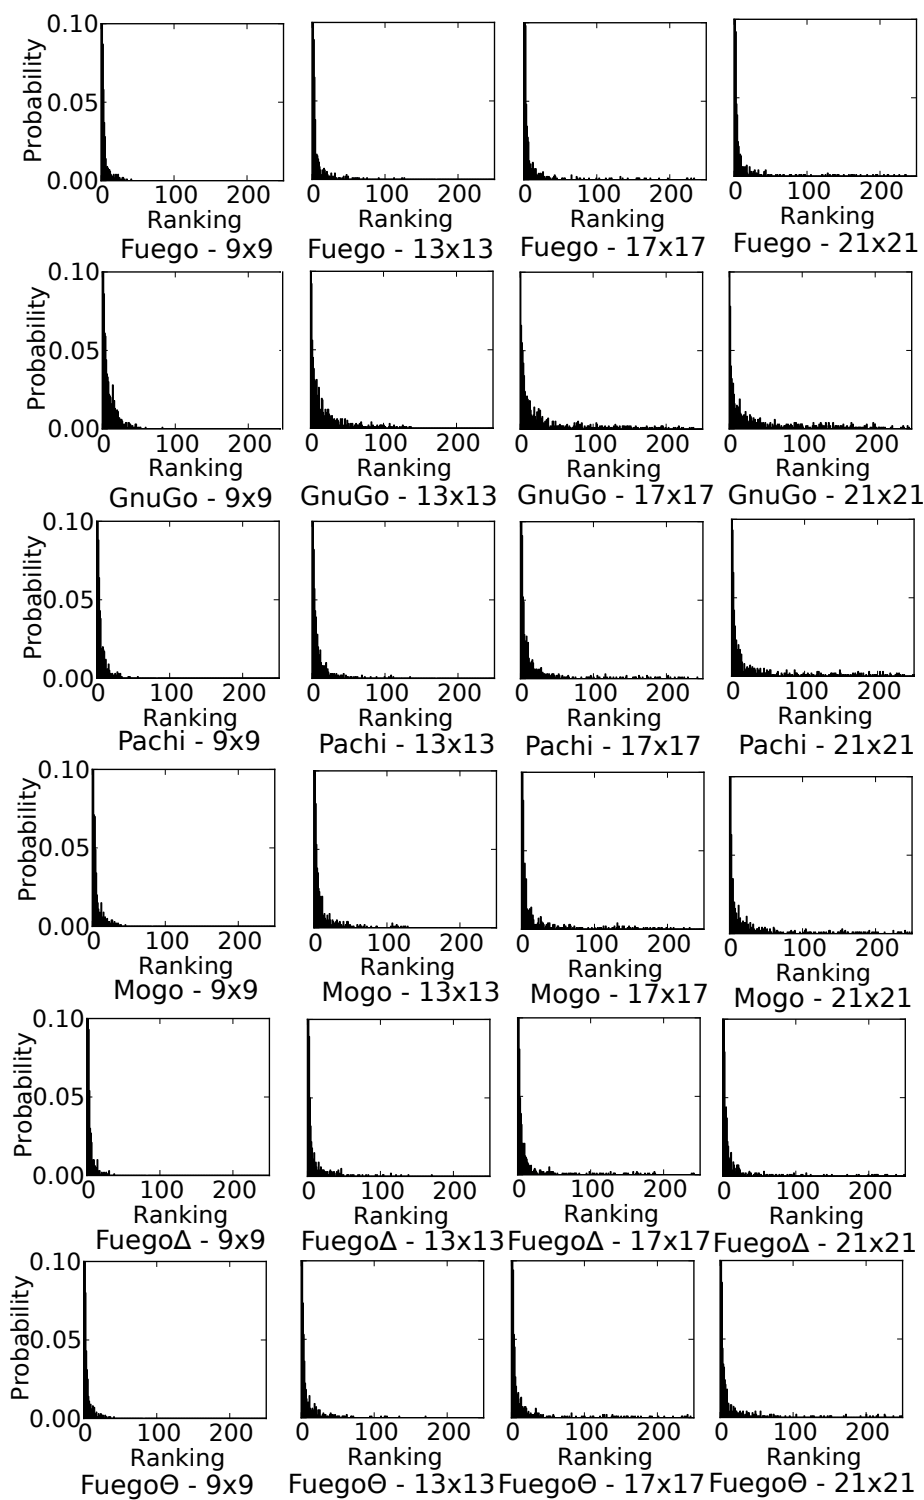


Fig. 5: Histograms of agents for different board sizes.

into bigger sets of actions as we increase the action space, and even though the strongest agent does not behave in the idealized way it does have a slower growth rate.

We also verify how the probability of playing the best action changes for each one of the agents as the number of actions increases. Figure 6 (b) shows that even though all agents experience a decrease in $p_{i,0}$, it does not decrease much. From 9x9, all the way to 21x21, we measure the following decrease: 20%, 23%, 39%, 26%, 28%, 22%, for Fuego, GnuGo, Pachi, Mogo, Fuego Δ and Fuego Θ , respectively. Hence, on average, they decreased about 25% from 9x9 to 21x21. Even though our assumption about $p_{i,0}$ does not hold perfectly, the predictions of our model are still verified. Therefore, the amount of decrease experienced is not enough to avoid that the diverse team increases in performance as the action space grows.

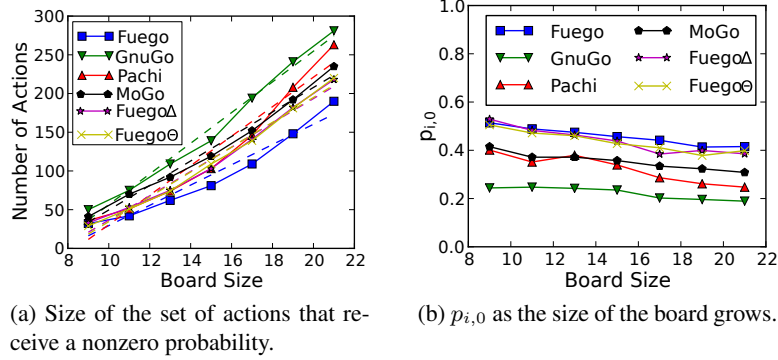


Fig. 6: Verifying the assumptions in the real system.

5 Conclusion and Discussions

Diversity is an important point to consider when forming teams. In this paper we present a new model that captures better than previous ones the intuitive notion of diverse agents as agents that tend to disagree. This model allows us to make new predictions. We show that the performance of diverse teams increases as the size of the action space gets larger. Uniform teams may also increase in performance, but at a slower pace than diverse teams. Therefore, even though a diverse team may start playing worse than a uniform team, it can eventually outperform the uniform team as the action space increases. Besides, we show that in large action spaces the performance of a diverse team converges exponentially fast to the optimal one as the number of agents increases.

We start our model with the notion of *spreading tail (ST)* and *non-spreading tail (NST)* agents. *ST* agents are agents that have a non-zero probability over a larger set of actions as the action space increases, while *NST* agents always have a constant number of actions with non-zero probability. We define a diverse team as a team of *ST* agents, and a uniform team as a team of *NST* agents. Therefore, our focus changes from modeling diverse teams as teams with different agents (as in models such as Marcolino, Jiang

and Tambe (2013) [14]), to focusing on diverse teams as teams where the agents tend to disagree. This change allows us to make new predictions that were not possible before.

Note that our model does *not* say that an *NST* agent will never vote for a new action. We define the pdfs of the agents by the rankings of the actions. Hence, when the number of actions increases from a certain number x_0 to a new number x_1 , a new action a^* may be the action with highest utility. Therefore, an agent will have the same probability of voting for a^* that it had for voting for the previously best action when the number of actions was only x_0 . A uniform team made of copies of the best agent also does *not* mean that the agents always vote for the same actions. The vote of each agent is a sample from a pdf, so copies of a single agent may or may not vote for the same action. In fact, we observe an increase in performance by voting among multiple copies of a single agent, both theoretically and experimentally.

The division of agents into two types (*ST* and *NST*) is, however, only an idealization, that allows us to isolate and study in detail the effect of diversity. A very strong agent will normally have most of its probability mass on the actions with the highest utility, so in the extreme its pdf would never change by adding new actions. In reality, however, it may also consider a larger set of actions as the action space grows. Therefore, we relax our model, and introduce the hypothesis that the best agent spreads the tail of its pdf at a slower pace than weaker agents. We show that because of this effect, a diverse team increases in performance faster than uniform teams, and we illustrate this phenomenon with synthetic experiments. Hence, even in a relaxed model where both diverse and uniform teams are composed of *ST* agents, a diverse team still outperforms a uniform team as the action space grows. The effect, however, is transient, as a uniform team may still have a higher convergence point than a diverse team, so in extreme large action spaces it would again outperform the diverse team. If the agents have the same probability of playing the best action, however, then it is clear that in the limit the diverse team will always be better than the uniform team.

Our model needs one strong assumption: that the probability of the individual agents voting for the best action does not change as the action space increases. This assumption allows our analysis to be cleaner, although it may not hold perfectly in a real system. In fact, in our Computer Go experiments we did observe a decrease in the probability of the agents voting for the best action. However, even though the assumption did not hold perfectly, the predictions of our theory holds: a diverse team significantly increased in performance as the action space got larger. Clearly, a decrease in the probability of the individual agents voting for the best action will decrease the performance of a team, while the effects studied in this paper will increase the performance. Therefore, as long as the decrease is not large enough to counter-balance the effect under study, we are still going to observe an increase in performance as the action space gets larger. Moreover, as we discuss in our generalizations, the argument that teams that spread the tail faster converge faster is still valid when the assumption does not hold; hence if the agents are equally strong (i.e., the individual agents have the same probability of voting for the best action) the team with faster growing tail will always perform better.

As mentioned, we verified our theory in a real system of Computer Go playing agents. Not only a real diverse team of agents effectively increased in performance as the board size increased, but we also verified that the strongest agent indeed spreads the

tail of its pdf at a slower rate than other weaker agents. We also verified that both diverse and uniform teams increase in performance, but the diverse team increased two times faster. This is explained by the relaxed version of our model, when we predict diverse teams to converge faster than uniform teams, as illustrated by our synthetic experiments.

Acknowledgments: This research was supported by MURI grant W911NF-11-1-0332.

References

1. Agmon, N., Stone, P.: Leading ad hoc agents in joint action settings with multiple teammates. In: AAMAS. pp. 341–348. Richland, SC (2012)
2. Barrett, S., Stone, P., Kraus, S., Rosenfeld, A.: Teamwork with limited knowledge of teammates. In: AAAI (2013)
3. Caragiannis, I., Procaccia, A.D., Shah, N.: When do noisy votes reveal the truth? In: EC. pp. 143–160 (2013)
4. Chen, H., Yao, X.: Regularized negative correlation learning for neural network ensembles. *IEEE Transactions on Neural Networks* 20(12) (2009)
5. Conitzer, V., Sandholm, T.: Common voting rules as maximum likelihood estimators. In: UAI (2005)
6. Fu, B., Wang, Z., Pan, R., Xu, G., Dolog, P.: An integrated pruning criterion for ensemble learning based on classification accuracy and diversity. In: Proceedings of the 7th International Conference on Knowledge Management in Organizations. KMO (2012)
7. Guttman, C.: Making allocations collectively: Iterative group decision making under uncertainty. In: Proceedings of the 6th German Conference on Multiagent System Technologies. vol. 5244, pp. 73–85 (2008)
8. Hong, L., Page, S.E.: Groups of diverse problem solvers can outperform groups of high-ability problem solvers. *Proceedings of the National Academy of Sciences of the United States of America* 101(46), 16385–16389 (2004)
9. Jiang, A.X., Marcolino, L.S., Procaccia, A.D., Sandholm, T., Shah, N., Tambe, M.: Diverse randomized agents vote to win. In: NIPS (2014)
10. Lamberson, P.J., Page, S.E.: Optimal forecasting groups. *Management Science* 58(4), 805–810 (2012)
11. LiCalzi, M., Surucu, O.: The power of diversity over large solution spaces. *Management Science* 58(7), 1408–1421 (Jul 2012)
12. Liemhetcharat, S., Veloso, M.: Modeling and learning synergy for team formation with heterogeneous agents. In: AAMAS. pp. 365–374. Richland, SC (2012)
13. List, C., Goodin, R.E.: Epistemic democracy: Generalizing the Condorcet Jury Theorem. *Journal of Political Philosophy* 9, 277–306 (2001)
14. Marcolino, L.S., Jiang, A.X., Tambe, M.: Multi-agent team formation: Diversity beats strength? In: IJCAI (2013)
15. Matthews, T., Ramchurn, S.D., Chalkiadakis, G.: Competing with humans at Fantasy Football: Team formation in large partially-observable domains. In: AAAI. pp. 1394–1400 (2012)
16. Nair, R., Tambe, M.: Hybrid BDI-POMDP framework for multiagent teaming. *Journal of Artificial Intelligence Research* 23(1), 367–420 (Apr 2005)
17. Polikar, R.: *Ensemble Machine Learning: Methods and Application*, chap. Ensemble Learning. Springer (2012)
18. Soufiani, H.A., Parkes, D.C., Xia, L.: Random utility theory for social choice. In: NIPS. pp. 126–134 (2012)
19. Thompson, A.: Does diversity trump ability? An example of the misuse of mathematics in the social sciences. *Notices of the AMS* 61(9) (2014)